

Management and Control of Transparent Optical Networks

Mari W. Maeda

(Invited Paper)

Abstract—Multiwavelength optical networking is expected to play a significant role in the next-generation transport networks providing capacity enhancements as well as built-in network survivability and reconfigurability. While advances have been made in the hardware technologies, considerable research effort is still required in the area of network management and control in order for optical networking to be proven commercially viable. This paper investigates key design issues concerning optical network management and control and examines how the networking architecture is influenced by the various management considerations.

1. INTRODUCTION

WAVELENGTH-DIVISION multiplexing (WDM) is widely becoming accepted as a technology for meeting growing bandwidth, and WDM systems are beginning to be deployed in both terrestrial and undersea telecommunications links. In today's commercial systems, wavelength multiplexing is used to enhance the fiber capacity strictly on a per-link basis with the wavelength connectivity remaining quasi-static; any reconfiguration requires manual rearrangement in the fiber patch panel or addition of new fiber connections or optical transmitters and receivers. Networking functions, such as aggregation and switching, remains in the arena of nodes capable of digital bit-processing. In contrast, research has been underway to explore the possibility of supporting networking functions directly in the optical layer, enabling a high-bandwidth (2.5 Gb/s and above) end-to-end circuit to be set up across optical crossconnect nodes via software control [1]–[4]. In addition to being dynamically reconfigurable, an optical network might be “transparent” with little or no bit-processing in the intermediate nodes. Fig. 1 shows an example of how high-speed transport layer signals might be routed over wavelength paths in an optical network. The basic functions supported by an optical network element (NE) include wavelength crossconnection, conversion, and multiplexing/demultiplexing.

While the term “transparent optical network” is widely accepted to imply dynamically reconfigurable network with no signal regeneration, “transparency” and “reconfigurability” are orthogonal attributes that can be supported to a relative degree.

Moving from static to software-controlled reconfigurable networks offers a number of advantages to both network providers and end-users. It enhances the efficiency of network usage through dynamic time-sharing of the resources, allowing end-to-end connections to be set up and released based on demand and traffic. Automatic provisioning opens up the possibility of enabling users to pay for high-bandwidth connectivities based on usage. With the development or enhancement of appropriate operations systems, the automated management features embedded into an optical network could dramatically reduce data inconsistency that plague today's networks and broadly reduce the operations costs through automation. Network survivability can also be efficiently addressed in the optical layer with the use of one of several possible restoration techniques [3]. A number of research programs are exploring the degree of dynamic reconfigurability that can be built into an optical network^{1,2} [1], [2].

While full optical transparency may be very difficult to achieve—especially with the need for all-optical wavelength conversion—a move to reduce regeneration and bit-processing is already seen in the commercial arena with an introduction of optical amplifiers that replace regenerators. Increased transparency in optical networks offers the potential for lowering the overall network cost by replacing electronic switches with simpler optical nodes. Cost-effective network evolution is made possible since the core optical hardware is independent of the transmitted signal format and rate, and future upgrades entail minimum equipment modifications. Transmission on a physically distinct wavelength also reduces the risk of signal tampering, since bit-level mixing (e.g., time-domain multiplexing or statistical packet multiplexing) or monitoring is not supported for the user traffic.

To date, researchers have been very successful in advancing optical technologies for reconfigurable optical networks. In order for a new technology to become commercially viable, however, network management and control functions need to be addressed as an integral part of the networking architecture. How is an end to end optical path set up? What mechanisms are supported to isolate and to diagnose fault? What NE to NE control functions are necessary for various survivable architectures? This paper explores the challenges of managing

Manuscript received October 6, 1997; revised April 1998.

The author was with Bellcore, Red Bank, NJ 07701 USA. She is now with the Defense Advance Research Projects Agency, Arlington, VA 22203-1714 USA (e-mail: mmaeda@darpa.mil).

Publisher Item Identifier S 0733-8716(98)05753-9.

¹ For information on WDM with Electronic Switching Technology (WEST) program, see <http://www.risc.rockwell.com/west>.

² For information on National Transparent Optical Network (NTON) program, see <http://www.ntonc.org>.




Fig. 1. Optical networking and client layer. The example shows two SONET regenerator section links being transported over optical networking layer. Link 1 (OC-12) signal is transported between SONET terminal equipment over optical networking nodes B-C-D-E on wavelength 1. Link 2 (OC-48) signal is routed over optical nodes A-C-D on wavelength 2.

and controlling optical networks that are both *transparent* and *reconfigurable*, with a focus on desired features that would enhance automation in operations flows.

The paper is organized as follows. Section II describes the optical networking architecture and details the NE hardware design. Section III describes configuration management functions including end-to-end connection setup and release procedures. Fault management issues and optical network survivability are investigated in Sections IV and V. Finally, issues concerning integration of optical network management with client layer management [e.g., asynchronous transfer mode (ATM), synchronous optical networking (SONET)] are described in Section VI.

II. TRANSPARENT OPTICAL NETWORKING ARCHITECTURE

A transparent optical network is composed of nodes capable of crossconnecting and multiplexing optical wavelengths [1]. Wavelength add-drop multiplexer (WADM) allows individual or subset of wavelengths to be added to or dropped from the fiber carrying multiwavelength traffic. A multiwavelength amplifier amplifies input WDM signals and differs from existing single-wavelength amplifier systems in its capability to monitor multiple wavelength and to perform fast gain adjustments in response to the dynamic changes in the number of wavelengths on the fiber. A wavelength-selective crossconnect (WSXC) can crossconnect an incoming wavelength to any outgoing fiber, while a wavelength-interchanging crossconnect (WIXC) has the added capability to perform wavelength conversion. At this time, most practical technique for performing wavelength conversion is through the use of a receiver/transmitter pair. These NE's might be physically placed in a combination of linear, ring, or mesh topologies depending on the demand pattern and the choice of architecture to support survivability.

Fig. 2 shows an example design for a WADM with single- and multiwavelength ports. A multiple-wavelength (east) input signal is demultiplexed and dropped via an optical crossconnect or transmitted out on the multiwavelength (west) port after multiplexing. Optical crossconnect switches may, for example, be based on electromechanic or electrooptic mechanism. The crossconnect shown in the figure must support internal signal terminations (or "dump" ports) for each wavelength so

that wavelength add action does not automatically result in wavelength drop action, and vice versa. An incoming single wavelength signal can be added to either an east or west multiwavelength output signal, and wavelength converters (in the form of receiver/transmitter pair) may be employed to assign the incoming signal to the desired channel. The wavelength assignment feature decouples the physical port from any particular wavelength and allows the network capacity to be more efficiently utilized.

Feedback controlled optical attenuators are used for channel equalization, i.e., to correct for nonuniformities in the incoming signal powers or amplifier gain spectrum. Optical amplifiers are used to compensate for link and component losses. In order to minimize cross-saturation effects and to stabilize per-channel gain, amplifier gain may also require feedback control [5]–[9]. In addition to the service-bearing (i.e., revenue-generating) wavelengths, an independent wavelength may be dedicated to support management and control channel (not shown in the figure). Numerous optical taps allow the integrity of the signal to be monitored at different points in the network elements.

It is instructive to compare optical networking with SONET in its overall architecture and goals. While SONET is a digital communications technology, the processing and manipulation of time-division multiplexed channels is quite analogous to the wavelength processing and multiplexing performed in optical networks [8], [9]. In SONET, lower-speed channels are multiplexed via byte-interleaving to generate a higher speed digital stream at up to 10 Gb/s. A SONET signal can simultaneously carry a mix of services at different digital rates, a feature analogous to optical network's goal to support different rates and formats. A winning feature that distinguishes SONET from old asynchronous technologies is its ability to directly extract or add specific channels and the associated payload bytes without having to process incoming signal in its entirety. This feature is analogous to transparent processing of optical channels and accounts for the relative simplicity and economy of optical networking nodes. In SONET, add-drop multiplexers (ADM's) and crossconnect nodes manipulate and provide access to individual channels, and analogous functions are supported in WADM's and crossconnect NE's. Time-slot interchange (TSI) and time slot assignment (TSA)

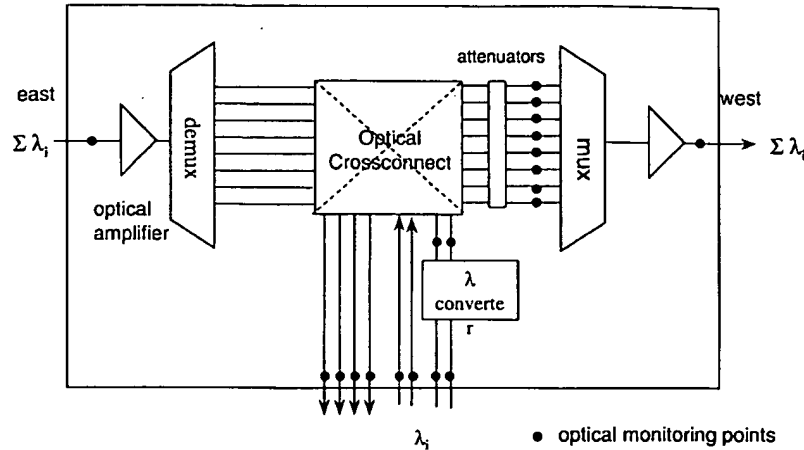


Fig. 2. Wavelength add-drop multiplexor. Only those components supporting signals in one direction are shown. Wavelength converter may consist of a receiver/transmitter pair.

features allow SONET payload to be flexibly assigned to a desired channel or time-slot whereas, in optical networks, channel assignment is done through wavelength interchange and assignment. Strong similarities are also evident between SONET self-healing architectures such as path-switched and line-switched rings and survivable optical ring architectures, as will be described in Section V.

While there are many similarities between SONET and optical networking, there are clearly many differences beyond the enormous bandwidth enhancement allowed with wavelength multiplexing. In the following subsections, we describe three key areas of distinction: 1) signal transparency; 2) absence of idle signals; and 3) the method to support management and control channel.

A. Optical Transparency

Optical transparency is often assumed to imply the possibility to transport signals of *any* modulation format (e.g., analog, digital) and data-rate across the network. While most optical components can be designed to be largely signal type independent, transmission limitations exist due to different *end-to-end* performance requirements for signals of different formats and rates [10]. Different signal types have different sensitivities to *cumulative* degradation from sources such as chromatic and polarization dispersion, amplified spontaneous noise, crosstalk noise, and optical nonlinearities. Also, it is difficult to support transmission of analog signals due to their sensitivity to optical reflections and due to stringent linearity requirements for the lasers used in wavelength converters. Steady-state behavior of the signal must also be considered (e.g., if the signal is not scrambled, a long stream of zeros could lead to average power drop) to ensure compatibility with the NE's signal monitoring capability. With careful engineering both in the hardware and the management software, however, it is still feasible to support a well-defined set of signal types such as digital signals at rates of up to 10 Gb/s. The implications of transparency to fault and configuration management functions are explored further in the next sections. The extent to which

signal regeneration, reshaping, and retiming may be desired will also be discussed.

B. Absence of Idle Signals

Another distinguishing feature of optical networks is the absence of idle signals at the service-bearing wavelengths. Prior to connection setup and service introduction, the only light that is propagated between NE's may consist of amplifier noise and the supervisory channel wavelength. Service-bearing wavelengths are introduced into the network only when a path is provisioned, hence, the actual number of wavelengths transmitted on a link at any given time may be fewer than the maximum number supported by the hardware, depending on the number of optical paths that have been configured. This contrasts with existing networks where optical carrier and overhead bytes are present independent of service, and any absence of optical signal or proper overhead bytes indicate a fault condition.

There are several alternative WDM networking architectures which support idle signals as shown in Fig. 3. In the first architecture [Fig. 3(a)], an electronic crossconnect fabric is used for crossconnection and the idle signals are always present. In the second architecture [Fig. 3(b)], optical crossconnect is used for switching but, again, the presence of optical transmitters for regeneration enable support for idle signals. Fig. 3(c) shows a node supporting transparent optical crossconnection, but it is also capable of dynamically inserting idle signals for channels that are not in use. An advantage associated with populating all wavelengths with light is the decoupling of fault and connection management functions, so that the absence of optical signal clearly signifies fault condition. Also, amplifiers used in various NE's might be operated under relatively stable conditions, relaxing the need for dynamic amplifier gain adjustments during service provisioning. Architectures supporting idle signals can be costly, however, requiring additional transmitter lasers. In the architecture illustrated in Fig. 3(c), network configuration also becomes very complex since a procedure for removing and adding idle signals must be supported by management systems or by NE's. And finally,

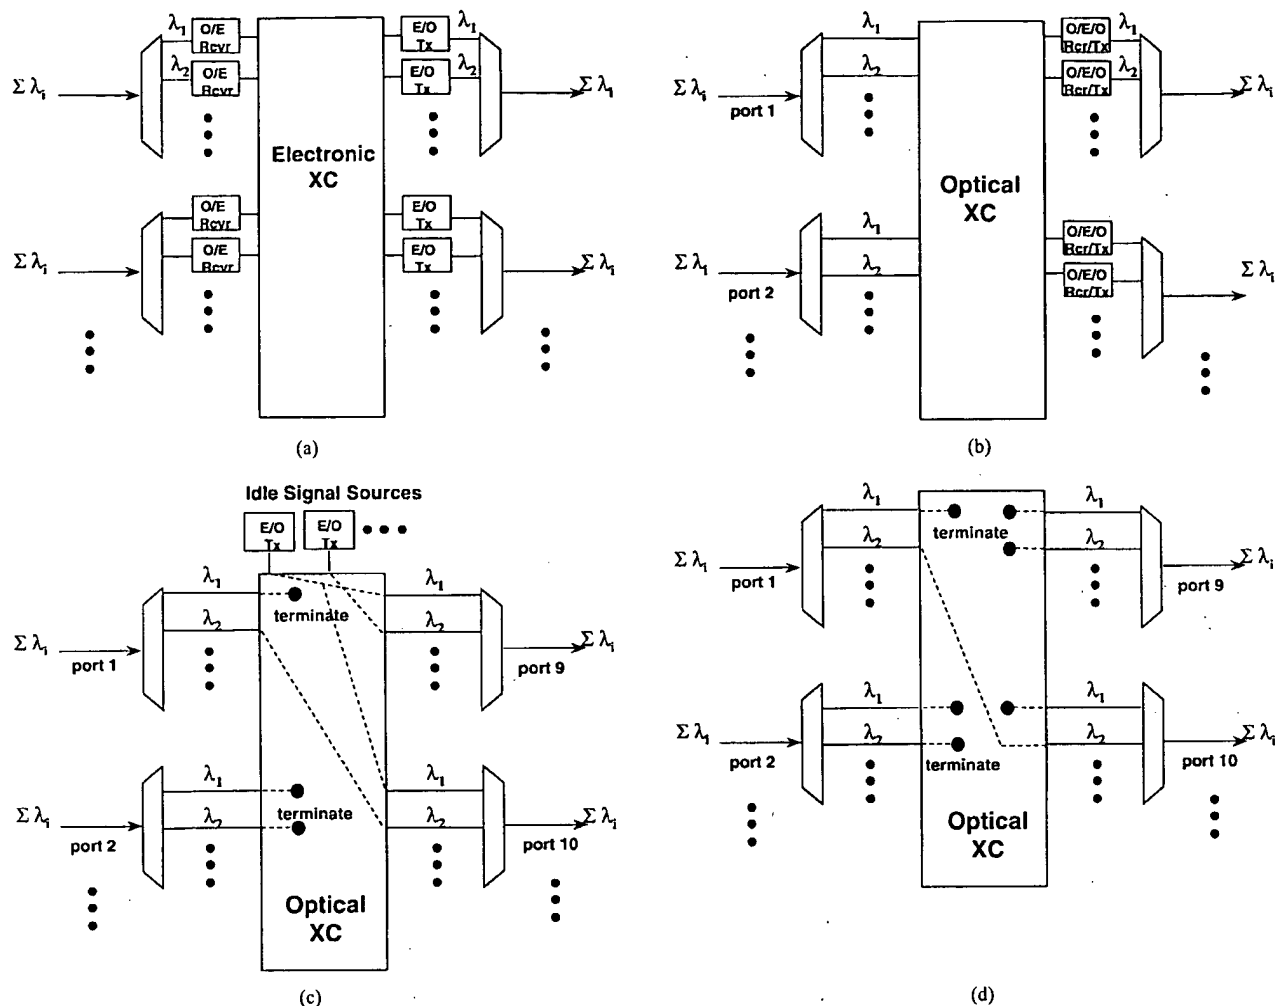


Fig. 3. Reconfigurable WDM crossconnect nodes supporting transparency and idle signals to varying degrees. (a) WDM node based on electronic crossconnect. (b) Optical crossconnect with signal regenerators. (c) Transparent wavelength selective crossconnect with idle signal transmitters. (d) Transparent wavelength selective crossconnect.

the seeming benefit of amplifier stability is not quite justified, since dynamic gain adjustment is still necessary to localize any impact of line failure. How the absence of idle signals impacts network management functions are explored further in Sections III and IV.

C. Optical Supervisory Channel

In existing transport systems, management channels are embedded in the digital data stream (e.g., ATM OAM cells and SONET overhead bytes). In the optical networks, because individual NE's do not have access to the bits transmitted over service-bearing wavelengths, a separate wavelength must be allocated to carry management and control information.³ Unlike the service-bearing channels which might travel transparently through a node, the optical supervisory channel (OSC) is electrooptically terminated, processed, and regenerated at

³An alternative method for transporting supervisory data may involve employing unused overhead bytes on the signals on bearer wavelengths; however such technique relies on assumptions about format of signals carried on the bearer channel and also limit utility at optical amplifier nodes.

each adjacent node. Whether there will be an OSC on every fiber between the multiply linked pair of nodes is a cost versus features tradeoff issue requiring further investigation.

An OSC will most likely transmit bit- and message-oriented data. A bit-oriented signal carries data that require fast real-time processing, such as indications of upstream failure or automatic protection switching signaling. It can only be used to carry a limited number of messages, however, and is quite limited in flexibility. On the other hand, message-oriented data require higher layer processing, but offers greater flexibility and can be used to support messaging between management and remotely located NE's for queries, software download, etc. Fig. 4 shows an example of a WSXC and how the OSC might be processed. Because the OSC wavelength may not fall within the amplifier gain bandwidth and because the transmission span is limited, OSC signals are demultiplexed and detected without propagation through an optical amplifier. After processing of the incoming management data, the outgoing OSC signal is generated and added to the appropriate output port. Some bit-oriented data—such as alarm indication signal and




Fig. 4. Wavelength selective crossconnect node supporting optical supervisory channel (OSC).

remote defect indication signal—are specific to a particular service-bearing wavelength, and are forwarded to the correct output port after processing. Thus, an OSC processor uses a routing table composed of input/output wavelength and port entries that mirrors the crossconnection state of the service-bearing wavelengths.

The specification of OSC content, its format, choices of signal rate, and wavelength must be standardized in order for multiwavelength mid-span meet to become a reality.

III. CONFIGURATION MANAGEMENT

Once the equipment and fibers are in place, software functions are invoked to configure the network. A network management system (NMS) constructs a view of the network by issuing queries to the NE's and monitoring notifications triggered by NE configuration updates. An NMS is also responsible for issuing requests to change the nodal configuration and to set up end-to-end optical paths [11], [19]. In the following subsections, we examine hardware and software requirements to support automatic topology discovery and connection setup.

A. Automatic Topology Discovery

A network's capability for self-discovery—to determine and communicate the current configuration to a managing entity—is of premier importance in simplifying the operations process. Removing the need for error-prone manual data entry promotes accurate and efficient information updates in the network database. An example is the NE self-inventorying function in which the NE automatically discovers the presence of bays, shelves, and plug-in units and determines their states; this feature is becoming widely supported in existing transport NE's. Automatic topology discovery is a feature of equal importance, enabling management systems to discover and verify link connectivities between adjacent NE's. New technologies, such as ATM switches and IP routers, are beginning to support automatic topology discovery, but are limited to discovering

logical connectivities in their specific networking layer. For example, an ATM switch might be able to identify its connectivity to an adjacent ATM switch, but cannot determine how the connecting trunk is routed over lower transport layers, e.g., over SONET multiplexing layer and physical medium layer. For an optical network, link connectivity discovery allows the management systems to build a physical medium (i.e., fiber) map and its discovery is of prime importance to subsequent optical path setup as well as for fault isolation. Here, three topology discovery methods are described (see Fig. 5).

1) *OSC-Enabled Adjacency Table Update:* In this method, an OSC is utilized to determine the adjacency of two NE's connected by the same fiber supporting the OSC. Because an OSC is electrooptically terminated at each adjacent node, an NE determines its link connectivity to the adjacent node and updates its adjacency tables. The method is illustrated in Fig. 5(a). NE-A queries NE-B through ports A.1 and B.1. Responses are returned via ports B.2 and A.2. Unique names or identifiers must be predefined for the NE's and the ports so that this information can be exchanged and used to form an entry into the adjacency table. The table entry should, at a minimum, include three fields: 1) port name; 2) adjacent NE name; and 3) adjacent NE port name. The fields may be automatically populated when links are brought in service or when prompted by the management system. The management system constructs a network fiber map by querying adjacency tables for all NE's. The advantage of this method is that the link connectivity can be verified even while the fiber is in service.

2) *OSC-Assisted Adjacency Table Update for Non-OSC Supporting Links:* Due to cost constraints, not every WDM link between the same pair of NE's may be supporting an OSC channel. If at least one pair of fiber to an adjacent node supports an OSC, however, it is possible for the nodes to determine the adjacency information over all the fibers. This method is illustrated in Fig. 5(b), and requires all ports to have optical sources [which could be in the form of laser

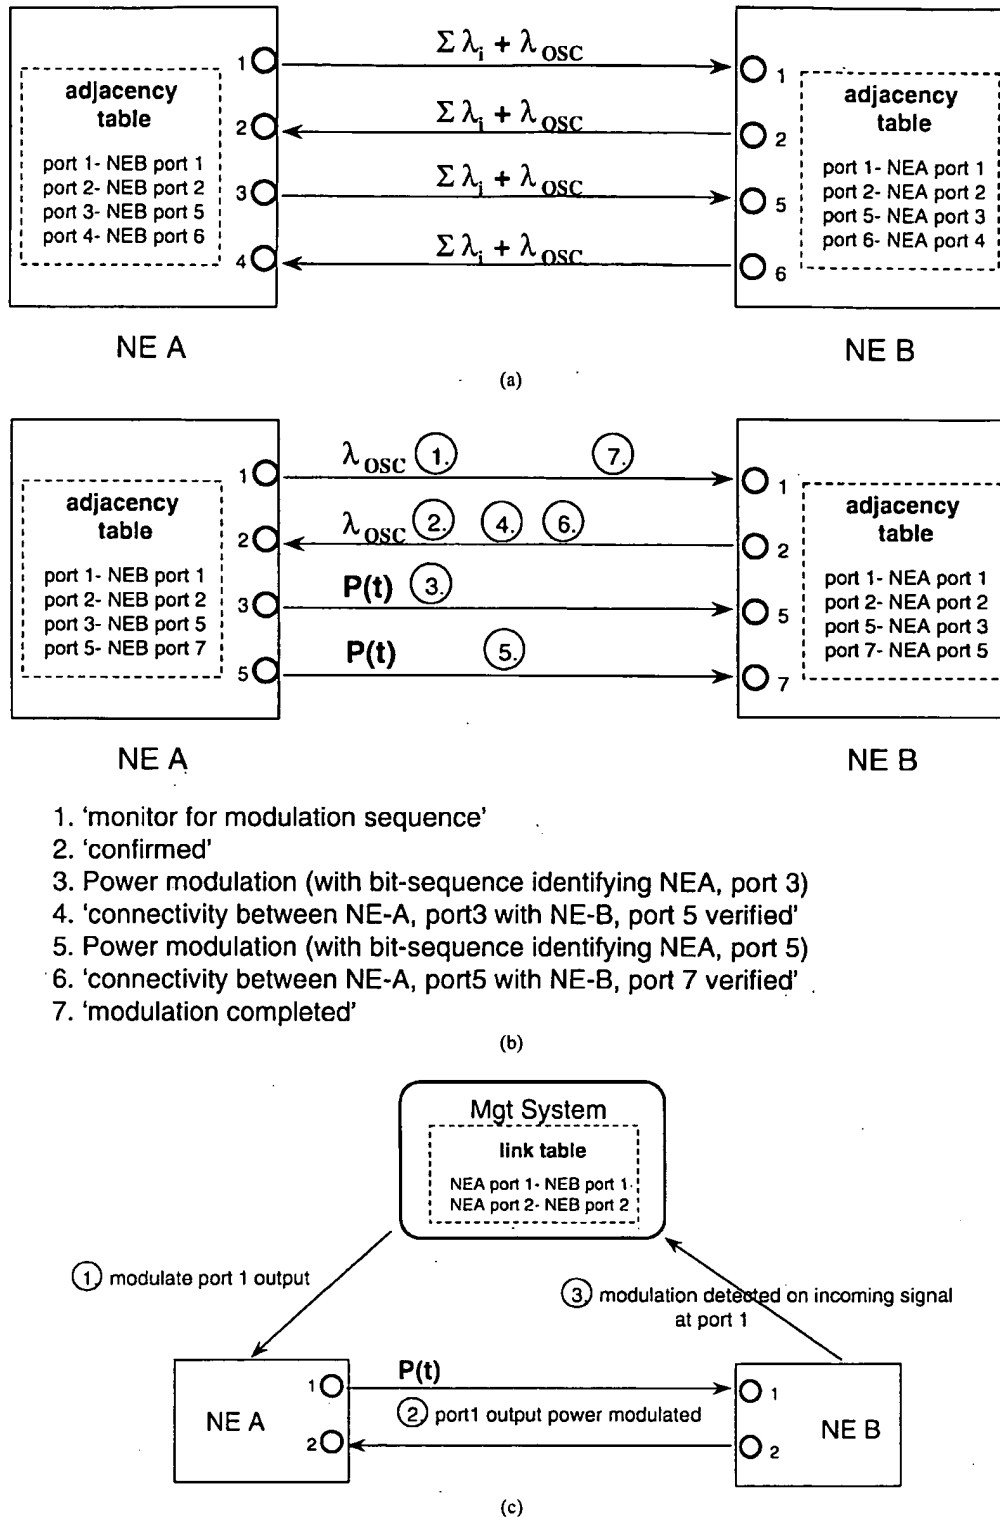


Fig. 5. Automatic fiber topology discovery. (a) Using OSC on every fiber. (b) OSC supported on subset of fibers. (c) Network management system-coordinated topology discovery.

transmitters or an optical amplifier with amplified spontaneous emission (ASE)] and optical sinks and modulation/detection capabilities. This method is invoked during the preprovisioning phase for the links in question.

Prior to introduction of service-bearing wavelength, NE-A registers with all adjacent nodes via OSC so that it will be notified of modulation detected on any of the incoming links. On receiving acknowledgment, NE-A proceeds to modulate

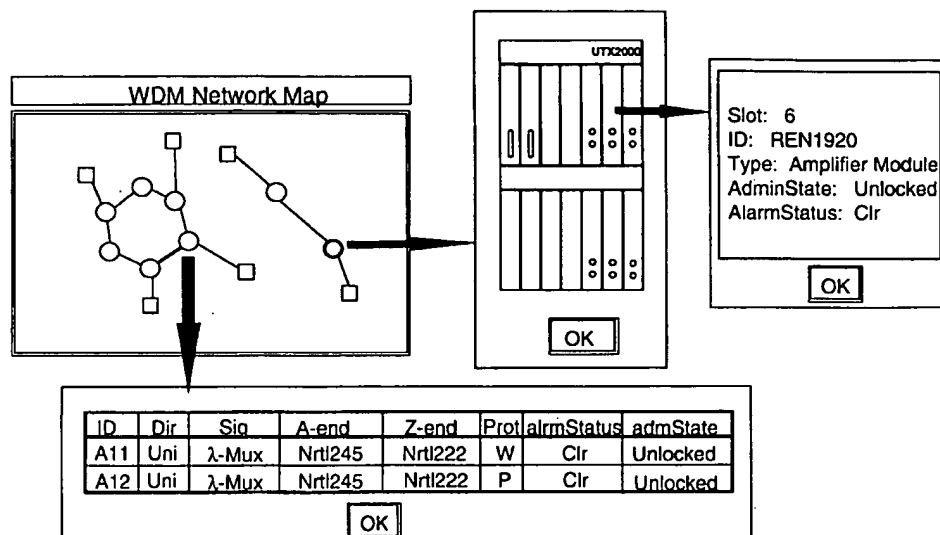


Fig. 6. Example of GUI supporting configuration management.

the output light (e.g., amplifier ASE) at the ports in question for a limited time interval with a predetermined bit-pattern. Via OSC, NE-B notifies the detected pattern to NE-A along with the NE-B port identifier, and NE-A uses the information to create or update an adjacency table identical to the one described in Method 1. Once the procedure is completed for all the links, NE-A deregisters itself from all the adjacent nodes. The drawback to this method is that the procedure can be invoked only during the pre-provisioning stage, and cannot be used once the service-bearing wavelengths are introduced. The simplicity of the scheme is quite attractive, however, requiring NE's to simply support signal sources and signal sinks.

3) Management System Coordinated Topology Discovery: With this method, the presence of OSC need not be assumed and an NE does not steward the adjacency table. However, an NE must be able to respond to a request from a management system to modulate the signal (or ASE) on any specified output port with a fixed bit-pattern, and to detect and report the presence of such modulation on any input port. Hence, by issuing such a request on successive NE-ports and monitoring for appropriate notifications, the management system determines the fiber links between NE ports. Again, this method requires link connectivity verification to take place prior to service provisioning so that each link is terminated or sinked and none of the signal wavelengths are transparently routed. The specifications to enable NE-to-NE interoperability, however, may be simpler.

Using information obtained during topology discovery, the management system may display a graphical user interface (GUI) representation of the network such as the one shown in Fig. 6. By clicking on the icon representing an NE, a network operator should be able to view the equipment faceplate and the results of NE self-inventorying. Also, a click on a line representing a link may bring up information such as the terminating port identifiers and whether the fiber is a single or multiple wavelength link.

B. Connection Setup

There are two styles for making an end-to-end optical connection: provisioning system based and control or signaling based. Provisioned connections require a management system application to coordinate route selection and crossconnect activation at NE's along the path, and these connections have a relatively long life-span. In contrast, switched circuit setup relies on peer-to-peer signaling between the control software loaded into NE's and is completed without management system intervention. The signaling could take place over an overlay network, similar to public switched telephone network today, or over OSC, similar to the way in which ATM switched virtual circuits are set up. Switched circuit setup is directly end-user initiated and fast. Signaling-based circuit setup requires careful standardization of the signaling protocol, however, if a mix of different vendor systems are to coexist within the network. While connection provisioning will be used in most near-term applications, signaling-based connection setup is being investigated as an advanced research topic [1].

1) Connection Provisioning: During connection provisioning, an optical path may be selected in one of several ways. The operator may hand-select the exact path for the connection, or may provide partial assistance, identifying one or a few intermediate points in the path. Alternatively, the route selection may be completely automated in which case a path provisioning system assigns a route based on input request parameters such as the identity of source and end points, source wavelength, and level of quality of service (QoS) required. A number of constraints must be met during transparent optical path setup and enforced with the aid of the management system, NE-specific software, or the equipment hardware.

- 1) Path provisioning must not impact existing services. An introduction or removal of a signal at any wavelength should not impact the signals on other wavelengths. Due to gain cross-saturation effects in optical ampli-

fiers, however, gain per channel is dependent on the total power propagated through the amplifier. Several feedback techniques have been proposed to maintain per channel gain constant. They are implemented in hardware—instead of NE or management system software—due to the speed requirement [6], [7].

- 2) An optical loop must not be created. This constraint is similar to one posed in existing digital technologies (e.g., IP), however an optical loop can have the dire consequence of immediately bringing down all the services in the path of the loop. A connection manager must verify that this constraint is met during route selection and during crossconnect activation procedure. Similarly, in order to avoid loop creation during protection switching (this topic will be discussed in Section V), NE software should trigger appropriate actions to drop any wavelengths for which a loop may be formed.
- 3) Hop count over an optical path must not exceed a maximum value over any given transparent span. (Note that this is in addition to the pre-engineering of nodal placement.) Because the signals of different rates and formats have different end-to-end performance requirements, a connection manager must be able to determine the maximum hop count, given the signal type, and choose a route that meets this constraint. The maximum hop count must also be satisfied even when the protection switching is activated. Alternatively, optical transparency may be limited to subnetworks (for example, optical rings) with regeneration at the intersubnetwork boundaries.
- 4) Appearance or disappearance of a wavelength due to provisioning and due to fault must be clearly distinguishable. Because idle optical signals are not supported, provisioning always results in appearance or disappearance of the carrier light. This physical activity must be clearly differentiated from fault condition by means of alarms or other mechanism built into the provisioning process. This issue is examined more closely in the next section.

2) *Optical Path Provisioning Sequence:* In existing high-speed networks based on optical transmission, the light over a particular link stays present and constant regardless of the provisioning process and an unexpected disappearance of light is an indication that there had been a failure. In optical networks that do not employ idle signals it is important to specify the signal propagation sequence during the provisioning process so that fault conditions can be clearly isolated.

Figs. 7 and 8 illustrate two procedures by which an optical path is provisioned. In both figures, points A and B represent client nodes and vertical lines signify intermediate NE's through which wavelength is propagated. A downward arrow indicates transition in time, thick diagonal arrows indicate light propagation followed by detection, and diamonds indicate the time and place at which crossconnect cut-through has taken place. Fig. 7 illustrates a connection setup algorithm in which the management system issues connection requests sequentially from upstream to downstream NE's. An NE activates a cut-through only if the expected incoming signal




Fig. 7. Connection setup procedure. Downward arrow indicates time and thick arrows indicate light propagation followed by optical detection. Half-diamonds indicate the time at which unidirectional crossconnect cut-through is activated. With sequential commit procedure, an NE activates cut-through only after verification of an incoming and outgoing wavelength signals. A connection request is rejected (i.e., an error is returned) if the incoming or outgoing signal is out of spec.

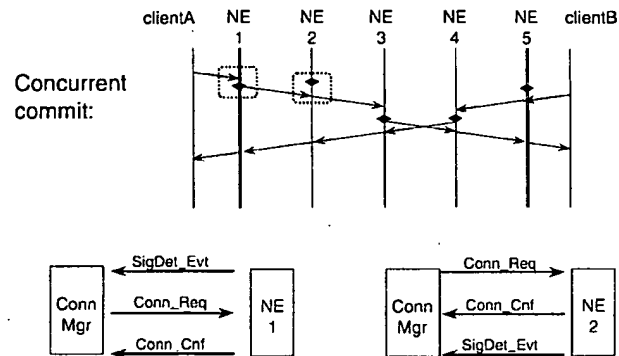


Fig. 8. Concurrent commit procedure involves simultaneous requests to all involved NE's to activate crossconnect cut-through. An NE returns success connection request confirmation after cut-through activation, without signal verification.

on a particular wavelength and port is observed (note that light detection always precedes an optical cut-through) and returns a successful confirmation to the management system on verifying the proper signal at the output port. On receiving the confirmation, the management system issues the next connection request to the following network element. This sequenced-verified connection setup method allows the management system to track the propagation of a new signal and promptly identify failures as indicated by detection of an unexpected wavelength or absence of an expected wavelength.

An alternate method is for the management system to issue concurrent connection requests to all the NE's without waiting for cut-through verification from each node. For some NE's (e.g., NE1 in Fig. 8) the incoming optical signal is detected prior to receiving the connect request, while for other NE's (e.g., NE2) a signal is detected only after receiving the connection request. Each NE is programmed to activate cut-through irrespective of the presence of incoming wavelength and to respond affirmatively once the crossconnection is made. With this routine, the management system relies

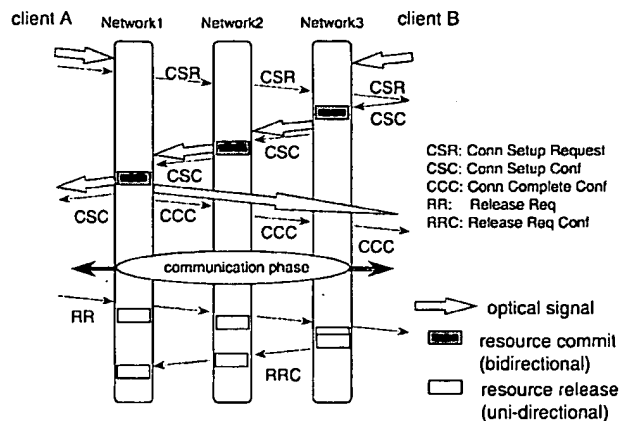


Fig. 9. Multidomain connection setup. In this example, an end-to-end connection is set up over networks managed by three carriers, whose management systems exchange messages to automate connection setup (see dotted arrows). Optical regeneration at network boundaries is not assumed, hence cut-through activation and optical signal propagation are closely coupled.

on separate event notifications from NE's—indicating that the appropriate signal has been detected—in order to verify successful end-to-end connection. If misrouting should be detected—i.e., notification indicates detection of signal at wrong NE or port—the management system must undo the previously committed cut-throughs, and invoke the diagnostic routine to isolate fault.

All connection release requests must flow sequentially from the optical sink to the optical source. This is necessary so as not to trigger alarms or restoration actions due to unexpected disappearance of a wavelength. For the connection shown in Fig. 7, the management system first issues a "disconnect" request to NE5, which deactivates cut-through, and returns a "disconnect" confirmation response. In the same manner, the management system sequentially issues "disconnect" request to NE4 all the way back to NE0. For optical networks, one must carefully define the fault monitoring architecture in conjunction with service provisioning architecture.

3) Multidomain Path Setup and Switched Circuit Setup: When an end-to-end connection spans multiple administrative domains—as would be likely in a national-scale networking environment—sequencing of light propagation and messaging will require further examination. Fig. 9 shows three networks, each domain with its own management system, and illustrates a possible procedure for setting up a transparent optical path from Client A to B. In this example, five main messages are exchanged between the management systems: 1) connection setup request (CSR); 2) connection setup confirmation (CSC); 3) connection complete confirmation (CCC); 4) release request (RR); 5) release request confirmation (RRC).

Network 1 receives a connection request from an end-user and forwards the CSR to network 2, which in turn forwards the CSR to network 3. The CSR message may contain the address of the end-users, the identification of the ports at the network boundary, and, possibly, the incoming wavelength. Once the end-network verifies the destination point and the end-user B accepts the connection request, network 3 activates a bidirectional connection so that the optical signal from client

B is propagated in the backward direction toward NE2 (only noise is propagated in the forward direction at this time). Network 3 also sends a CSC message back to network 2 which activates an optical cut-through. Finally, when all the connections have been activated in all three networks, the optical signal in the forward direction—from network 1 to 3—is transmitted and the CCC messages are issued so that the two end-users can begin communicating. This example of multidomain path setup protocol must be carefully enhanced with error handling routines.

Another management consideration in designing multidomain optical networks includes coordination of wavelength usage and impairment assignment. The difficulty of assigning end-to-end wavelength(s) based on the availability in multiple networks will most surely necessitate placement of wavelength interchanging crossconnect NE's at the network boundaries. Similarly, the enforcement of end-to-end signal quality by, for example, limiting the hop count, may prove too difficult to coordinate across multiple domains. Fault isolation and testing is also a big challenge. While optical transmission over a national-scale network is technically feasible [12], management considerations argue for limiting transparency at domain boundaries, and placing NE's that allow wavelength reassignment and signal regeneration.

4) Switched Connection Setup: An optically switched circuit network gives end-users direct control of the network resource, relying on NE control software and signaling instead of a management system, to set up end-to-end connections. Fig. 9 illustrates the signaling and signal propagation process if three networks are replaced by three NE's, and the arrows indicate messaging between NE control software. The signaling protocol and the signal propagation procedure must be carefully designed, taking into account various exception-handling procedures. Because the holding times of connections can be arbitrarily short for switched circuits, a management system cannot be expected to track creation and release of connections in real time. This fact introduces a new challenge to performing fault diagnostics, and a procedure for releasing faulty circuits and performing optical tests must be carefully analyzed. Similarly, when the ability to carry switched circuits is combined with the ability to carry different signal types, many new management challenges are introduced.

IV. FAULT MANAGEMENT

Fault management consists of maintenance functions such as network surveillance, and testing, as well as restoration [11]. NE-specific fault management features include: 1) monitoring and detecting fault conditions; 2) correlating internal and external symptoms; 3) reporting alarms to the management system; and 4) configuring restoration mechanisms. Network level fault discovery is made through alarm surveillance, detection of secondary symptoms (e.g., performance trends), or through customer call-in. The fault manager must isolate the source of a problem through alarm analysis and diagnostic tests so that the smallest replaceable equipment or cable span can be identified. In order to correct the problem, the fault manager interfaces with a trouble ticketing or dispatch

system so that appropriate repairs can be made. In this section, fault management functions of particular relevance to optical networks are reviewed.

A. Alarm Surveillance

Alarm notifications are autonomous messages generated by NE's indicating presence of a fault condition. Examples include: 1) plug-in equipment alarms (e.g., amplifier pump card failure, receiver/transmitter card failure); 2) environmental alarms (e.g., fire detected); and 3) software failure alarms. Alarms of particular relevance to optical networking are communication alarms indicating signal loss or degradation, and are examined in this section.

Fig. 2 shows examples of signal monitoring points within an NE. The monitored parameters include: 1) optical signal power; 2) optical noise levels [or signal to noise ratio (SNR)]; and 3) wavelength registration. Monitoring points must be placed to enable NE's to differentiate external problems from internal problems, and to isolate any internal fault. Certain tap points are used to derive signals used for feedback adjustments; two examples are the channel equalization made possible through individual channel monitoring and amplifier gain adjustment through total power monitoring. When a signal cannot be adjusted to fall within an acceptable range, an alarm notification is issued. Alarming should be software-configurable and alarm thresholds for most parameters should eventually be standardized. Because the network can simultaneously carry signals at different rates, per-channel thresholds such as SNR limits may have to be configured on a per-connection basis. For example, the SNR requirement is much more stringent for a wavelength carrying 10-Gb/s signal compared to a channel carrying 600 Mb/s. Hence, during connection setup, threshold configuration requests may need to be issued to individual NE's, together with connection setup requests. An alternative is to use a common threshold value for all channels, and to select this threshold to meet the most stringent signal requirement. A drawback of such an approach is that an alarm may be generated even when the optical signal degradation does not impact the digital service. Reference [13] describes yet another scheme where the alarm surveillance between the optical and digital layers are coupled via supervisory channel communications. Thresholds are set for the optical signals but the actual alarms are generated only if optical signal degradation is accompanied by digital signal failure (e.g., loss of frame). Such a method requires interoperable signaling between optical NE's and client digital NE's, and is feasible only with a redesign of client layer NE's.

Because transparent optical networking architecture does not support digital signal monitoring, one very real concern is that there may be failure modes that are difficult or perhaps impossible to isolate by means of optical monitoring alone. One example is an optical crossconnect failure that results in a placement of correct wavelengths at the right ports, however with the incorrect digital information content. Another example is a failure in the electronics of the wavelength interchanging module (receiver/transmitter pair), so that jitter or other corruption in the modulation signal is introduced. Once

the problem is detected (most likely through customer call-in or an alarm in the client layer NE) fault isolation will require testing beyond simple optical monitoring. While there may be solutions to some of these problems—for example, hardening or providing local monitoring to the individual components, or introducing pilot tones to track individual optical paths [14]—a practical approach is to limit the extent of transparency to smaller subnetworks and supporting some basic electronic monitoring and signal regeneration at the edges.

B. Alarm Propagation

Alarm propagation is another fault-related issue unique to optical networks due to lack of idle signals. Because optical signals are transparently routed through nodes, the same degradation can be observed at all downstream nodes. Using an effective alarm coordination mechanism, redundant alarms are generated by all the NE's along the path. For a small optical network or a subnetwork of limited optical transparency, the management system may receive, analyze, and suppress all the redundant alarms. For larger transparent networks, however, a single fiber cut, impacting, for example, 16 wavelengths, could trigger over 100 alarms. Hence, it is necessary to examine inter-NE communication mechanism by which redundant alarms are automatically suppressed [15]. Fig. 10 illustrates an alarm coordination mechanism using an alarm indication signal (AIS), a maintenance signal used to alert downstream node of an upstream failure. For optical networks, an AIS is transmitted over the optical supervisory channel, possibly by defining bytes that carry alarm status information for individual wavelengths. An NE detecting an incoming AIS determines that the source of failure lies upstream and suppresses local alarms.

For simplicity, Fig. 10 shows an impact of failure on an optical path 1 (λ_1) originating at WTM1 and terminated at WTM4. An assumption is that optical paths had been set up using the "verified sequenced setup method" described in Section III-B such that an absence of wavelength in the presence of a cut-through is an indication of failure. When the fiber is cut, WSXC1 detects total loss of incoming signal (including ASE noise) and begins transmitting an AIS downstream for all the wavelengths for which local cutthrough had been activated. The AIS signal may indicate the identification of NE detecting failure (i.e., WSXC1) and the type of failure for different wavelengths (e.g., loss of signal, wavelength drift). WSXC2 and WSXC3 also begin transmitting AIS along the same path for λ_1 , and when they detect an incoming AIS, they overwrite the AIS which they were transmitting downstream. WSXC2 and WSXC3 suppress alarms and WTM4 may issue an alarm indicating AIS detection. In this way, the total number of alarms is reduced to the primary alarm indicating total loss of signal at WSXC2, and the secondary alarms indicating AIS detection at the terminating node for all impacted optical paths (WTM4, etc.). In order to alert upstream nodes of downstream failure, remote defect indication (RDI) signals may also be defined. RDI's would similarly be transmitted over an OSC and may be instrumental in triggering upstream optical path protection switching.

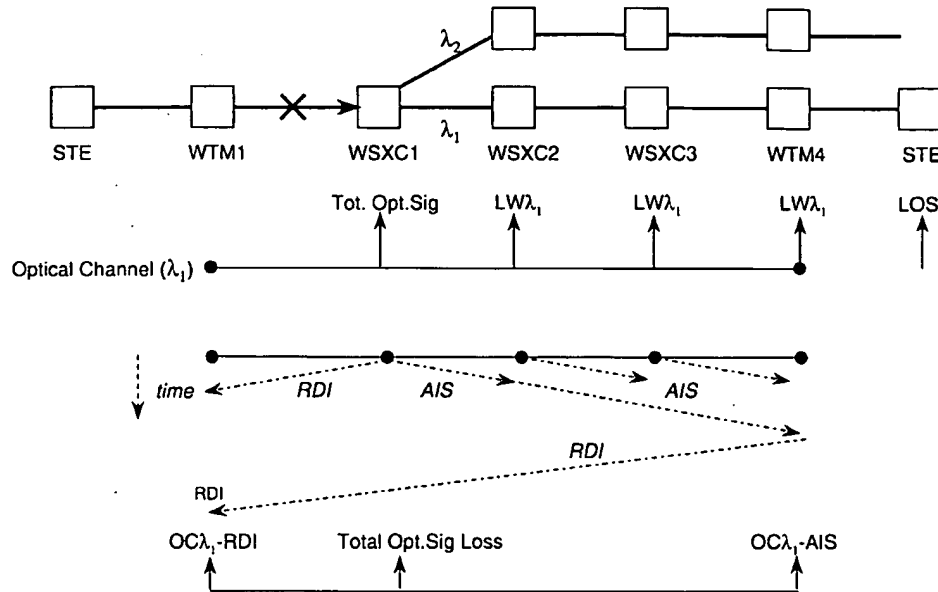


Fig. 10. Redundant alarm management scheme. OSC is used to transmit remote defect indicators (RDI's) and alarm indication signals (AIS's) along the wavelength paths. Vertical arrows indicate alarms generated. Top: without alarm coordination mechanism. Bottom: using AIS to suppress redundant alarms. LW: Loss of wavelength alarm. LOS: loss of signal alarm.

In certain situations, it is desirable to simply disable per-channel signal monitoring at intermediate nodes. Arguments for such arrangement are: 1) complex AIS mechanism becomes unnecessary and 2) it may not be desirable to monitor power level of certain signals that may be carrying optically bursty traffic or have characteristic such that average output power is not continuous and constant. One may also argue that any alarms associated with individual wavelength are secondary in nature, with primary alarms being due to equipment failure or fiber cuts. A sound alarm surveillance policy, however, is to produce more than one alarm notification per fault because of situations where primary alarms never reach the management system due to, for example, node failure or data communication network failure. With the AIS transmission technique described above, the network guarantees that more than one alarm is generated for any failure that impacts optical paths.

V. OPTICAL NETWORK SURVIVABILITY

Building network survivability into an optical layer offers many benefits [16]–[18]. Optical layer restoration can be efficient and cost-effective, requiring fewer overall hardware components and signal ports compared to client layer restoration. The restoration response speeds can be enhanced with the reduction in the number of restoration control messages and the use of a high-speed optical supervisory channel. While the most efficient use of resources is made with restoration architectures that rely on real-time rerouting algorithms, automatic protection switching (APS) rings with dedicated spare capacity are often preferred due to their superior restoration speeds. Examples of optical APS rings and their requirements for control are examined in this section.




Fig. 11. Wavelength path switched ring. Protection switching operates on individual wavelengths.

A. Optical Path Protection Rings

Fig. 11 depicts a survivable optical ring architecture that protects traffic on a per-channel (i.e., per-wavelength) basis. An optical path protection ring is analogous to a SONET unidirectional path-switched ring, and the ring nodes are connected with two sets of fibers carrying counterpropagating signals. Each wavelength that is added onto the ring is divided with an optical splitter and transmitted in both clockwise and counterclockwise directions. The return path is likewise split over two directions. The wavelength-dropping NE monitors both incoming signals and selects one with superior signal quality, hence, no signaling between NE's is required to coordinate protection switching. Also, because each wavelength is protected independently, the service is assured even for failures that impact a single wavelength. Each connection takes up the entire ring capacity for one wavelength, however, so the maximum number of optical paths that can be supported by a ring is equal to the total number of wavelength channels available.




Fig. 12. Two-fiber shared protection optical ring. Protection capacity is shared between working channels. The working channels on the inner ring are protection channels on the outer ring. Fiber cut triggers loopback switching on adjacent nodes.

B. Shared Protection Ring

In the two-fiber shared protection optical ring architecture (see Fig. 12), half the ring capacity is reserved for protection traffic just as it is in the path switched architecture, however, higher overall capacity can be attained because the spare capacity is shared between many connections. This architecture is similar to a SONET two-fiber bidirectional line switched ring, however due to the expense of implementing channel reassignment (i.e., wavelength-interexchange), different wavelengths are allocated for working and protection traffic in the two counterrotating directions. For the example shown in Fig. 12, λ_5 is used for an optical path from NE1 to NE4, whereas λ_1 is used for the return path from NE4 and NE1. A link failure activates loopback on multiplexed signal at the adjacent nodes so that the working traffic on the outer ring is routed on the inner ring, and vice-versa. NE-to-NE signaling is necessary to coordinate protection of switch activation and switch reversion.

Fig. 13 shows four-fiber protection shared ring architecture with two fibers carrying working traffic and two fibers carrying protection traffic. Each NE supports optical switching to transfer signals from working to protection fiber. This architecture is analogous to a SONET four-fiber bidirectional line switched ring and, like SONET, supports two different modes of switching depending on the failure modes. If only working fiber(s) fail, *span switch* is activated, and the multiplexed traffic between adjacent nodes (e.g., NE2 and NE3) is routed over the protection fiber over the same span [see Fig. 13(b)]. Span switching is a very useful feature for scheduled maintenance and testing purposes since its impact remains local, and it is often cited as the reason for the popularity of four-fiber rings. If both working and protection fibers are cut, *ring switch* goes into effect, and the failed link is replaced by a longer protection route around the ring [see Fig. 13(c)]. For example, if a cable cut occurs between NE2 and NE3, the entire multiplexed traffic from NE2 to NE3 is routed in the counterclockwise direction over the protection ring. Hence, an optical path from NE1 to NE4 is routed in the following sequence along the ring: NE1-NE2-NE1-NE6-NE5-NE4-NE3-NE4. The ring must be designed so that the end-to-end performance constraint is met for all optical paths in the protection-switched state. NE-




Fig. 13. Four-fiber shared protection optical ring. (a) A pair of fibers is dedicated for protection traffic. When fault is not present, traffic is routed over a working pair of fibers (see connection between nodes 1 and 4). (b) If protection fibers are intact, failure triggers span switch at two adjacent nodes (nodes 2 and 3). (c) When failure takes down all four fibers, ring switch is activated (nodes 2 and 3).

to-NE signaling is also necessary in order to activate and revert ring and span switching.

The signaling requirement must also be made in consideration of node failure and multiple failure scenarios. Fig. 14(a) and (b) depicts an example of a node failure. Node failure triggers a ring switch in the two adjacent nodes (i.e., NE2 and NE6) so that route NE6-NE1-NE2 is replaced by the longer route NE6-NE5-NE4-NE3-NE2. During this switching, care must be taken to remove or squelch all the connections terminating at the failed node to avoid misconnections. As




Fig. 14. Four-fiber shared protection optical ring and node failure. (a) Two connections between node 5 and node 1, and between node 1 and node 4. For simplicity, only two of four fibers are shown. (b) Node failure triggers loopback switching to take place at nodes 2 and 6. Without proper signal squelching, a misconnection occurs between node 4 and node 5.

shown in Fig. 14(b), without signal squelching, two independent connections between NE1 and NE4, and between NE5 and NE1 that utilize the same wavelength collapse into a single connection between NE5 and NE4, with the result of misrouting. Hence, signaling must be accompanied with appropriate actions to ensure that the NE's supporting connections to the failed NE release their crossconnections.

Fig. 15 illustrates further signal squelching considerations involving multiple failure scenarios. In this example, the first failure occurs over link NE1-NE6, followed by a second failure over link NE2-NE3. Ring switching is activated at NE1 and NE6 as a result of the first link failure. Following the second link failure, a second set of protection switching is activated at NE2 and NE3. At this point the ring becomes segmented into subnetworks (NE1 and NE2) and (NE3, NE4, NE5, and NE6), and the NE's must take care to squelch or release connections spanning across the two segmented subnetworks. This additional squelching is necessary to prevent the remaining crossconnections from




Fig. 15. Four-fiber shared protection optical ring and multiple failure scenario. (a) A connection between nodes 5 and 4 is routed over a long route due to link failure between node 1 and 6. (b) Second failure occurs between node 2 and 3, and second set of loopback switching is activated. Without proper signal squelching, an optical loop is created over one of the isolated subnetwork, potentially impacting all connections between node 1 and 2.

forming an optical loop, as shown in Fig. 15(b) between NE's 1 and 2. With a near unity net gain, an optical loop can cause service failure for *all* wavelengths sharing the path; even those connections confined to the segmented subnetwork would be impacted if squelching is not properly supported. Hence, proper termination of connections spanning segmented subnetworks must be coordinated between the NE's at the time of second protection switch.

This section reviewed several survival optical ring architectures and their control requirements. The choice of preferred protection architecture will depend on a number of issues, including the effectiveness of restoration, the efficiency of capacity usage, and the complexity of control. Integrating optical layer restoration to complement survivability mechanisms in the client layer is also an extremely important considera-




Fig. 16. Integration of multilayered network. (a) Physical connectivity between WDM, SONET, and ATM NE's is shown. (b) SONET and ATM element management systems (EMS) manage the network based on layer-specific logical connectivity of nodes.

tion and will influence optical survivability architecture and, possibly, impact client layer mechanisms already in place.

VI. INTEGRATED NETWORK CONTROL AND MANAGEMENT

The *layered* transport architecture offers many advantages, including the possibility for each network layer to evolve independently of the other layers. Care must still be taken to ensure interoperability of the various layers, however, especially in their management and control functions. The challenge is twofold. First, NE-to-NE interoperability features in both hardware and software need to be identified. For example, new standards for an optical and client layer NE may be necessary to ensure coordination of optical layer and client layer restoration or alarm generation mechanism. Second, an NMS that serves to integrate layer- and vendor-specific element management systems (EMS) will need to be designed and developed. Such an integrated management system (IMS) will provide the network operators with a common GUI to monitor and control the different layers of the network in a seamless fashion. Some of the key areas of concern for integrated management are described below.

1) *Integrated Topology Management*: Integrating configuration management across several networking layers requires an IMS that can tie together information provided by a layer-specific EMS, and definition of key NE-to-NE functions that are not supported today. Fig. 16 illustrates an example of three networking layers consisting of ATM, SONET, and optical network elements. For example, an OC3c signal from an ATM switch is multiplexed onto OC48 signal at a SONET multiplexer and transported over λ_1 in the optical layer.

While the logical networks managed by layer-specific EMS are viewed as shown in Fig. 16(b), an IMS ties the three views together and is able to display the physical connectivity of the various network elements as shown in Fig. 16(a). Certain information can be obtained from simple communication with an EMS system. For example, in order to view the state of equipment or ports, an IMS issues a query to the EMS under whose jurisdiction the NE falls. Information associated with links spanning NE's in different layers (e.g., a fiber link between a SONET NE and an optical NE), however, must be entered via the IMS and be verified via queries to separate EMS's. A greater challenge is for the NE's to directly support interlayer auto-topology discovery such that, for example, a SONET NE and an optical NE communicate with one another to determine their link connectivities. Such a feature would allow IMS to directly determine the network topology without relying on manually entered configuration data.

2) *Route Correlations*: Determination of the relationship between the client-layer links and a server-layer route is another key function for an IMS. Without an IMS, an ATM-layer EMS has no way of determining how an ATM trunk is routed over SONET multiplexers or optical layers. IMS must coordinate interlayer routing so that a network operator can select a physically diverse route for higher layer services. IMS can also facilitate grooming so that the fill efficiency can be enhanced for the lower layer capacity.

3) *Integrated Fault Management*: Coordination of root cause analysis across multiple technology layers is essential for improving operations efficiency. In a separately managed network, one failure in the optical layer would trigger a huge number of redundant alarms from the terminating nodes in

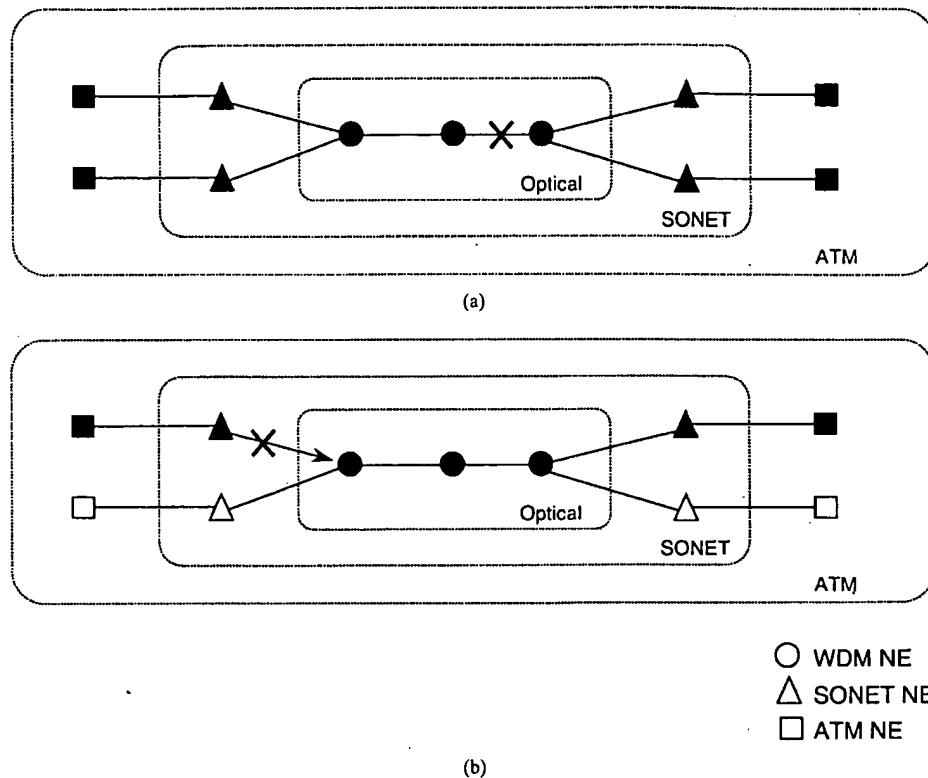


Fig. 17. Redundant alarms generated across multiple layers due to (a) fiber cut and (b) SONET transmitter failure. NE's detecting and reporting failure are darkened.

the client layers as illustrated in Fig. 17. The IMS needs to establish the dependencies between the various equipment, links, and connections across the layers and suppress redundant alarms across multiple networking layers. The IMS must also coordinate testing and setting of alarm thresholds based on the needs of the client layers.

VII. SUMMARY

Transparent optical networking is based on two basic ideas: optical transparency (minimal regeneration or bit-processing) and software-control reconfigurability. This paper investigated various network management and control issues for a transparent optical network where signal regeneration is minimized or not permitted at the intermediate nodes. Some of the challenges that arise for transparent networks include problems of fault isolation, alarm redundancy management, and cut-through timing coordination; each of which may not be insurmountable but would inevitably require rigorous standardization and complex software systems to be designed, especially if we hope to have a network that is interoperable across both multiple supplier systems and multiple administrative domains. The extent of complexity in management and control is greatly reduced when regenerative nodes are placed to confine optical transparency to smaller subnetworks or to a single administrative domain.

Built-in reconfigurability for high-speed optical paths offers many exciting possibilities and needs to be pursued actively as a research topic. The extent of transparency that is de-

sired in such reconfigurable networks, however, needs careful investigation, especially from the perspective of network management.

ACKNOWLEDGMENT

The author gratefully acknowledges stimulating discussions with MONET participants and support from DARPA.

REFERENCES

- [1] R. E. Wagner *et al.*, "MONET: Multiwavelength optical networking," *J. Lightwave Technol.*, vol. 14, pp. 1349–1355, June 1996.
- [2] I. P. Kaminow, "A wideband all-optical WDM network," *IEEE J. Selected Areas Commun.*, vol. 14, pp. 780–799, June 1996.
- [3] K. Sato, *Advances in Transport Network Technologies*, 1996.
- [4] S. Johansson, "Transport network involving a reconfigurable WDM network layer—A European demonstration," *J. Lightwave Technol.*, vol. 14, pp. 1341–1348, June 1996.
- [5] J. Zyskind *et al.*, "Fast power transients in optically amplified multiwavelength optical networks," presented at OFC'96, Postdeadline Paper PD31.
- [6] J. L. Jackel and D. Richards, "All-optical stabilization of cascaded multichannel erbium-doped fiber amplifiers with changing number of channels," presented at OFC'97, Dallas, TX, Paper TuP4.
- [7] J. L. Zyskind *et al.*, "Fast link control protection for surviving channels in multiwavelength optical networks," presented at ECOC'96, Oslo, Paper ThC.3.6.
- [8] M. Sexton and A. Reid, *Transmission Networking: SONET and the Synchronous Digital Hierarchy*. Norwood, MA: Artech House, 1992.
- [9] "Synchronous optical network (SONET) transport systems: Common generic criteria," Issue 1, Bellcore GR-253-CORE, Dec. 1994.
- [10] L. Gillner, "Transmission limitations in the all-optical networks," presented at ECOC'96, Oslo, Norway.
- [11] A. Leinwand and K. Fang, *Network Management*. Reading, MA: Addison-Wesley, 1993.

- [12] R. S. Vodhanel *et al.*, "National-scale WDM networking demonstration by the MONET consortium," presented at OFC'97, Dallas, TX, Postdeadline Paper PD-27.
- [13] Y. Tada *et al.*, "OA&M framework for multiwavelength photonic transport networks," *IEEE J. Select. Areas Commun.*, vol. 14, p. 914, June 1996.
- [14] F. Heisman, "Signal tracking and performance monitoring in multi-wavelength optical networks," presented at ECOC'96, Oslo, Norway.
- [15] M. Huber and O. Jahreis, "Supervision and protection concepts for an optical network," presented at OFC'95, Paper WO5.
- [16] T. H. Wu, *Fiber Network Service Survivability*. Norwood, MA: Artech House, 1992.
- [17] K. Oda *et al.*, "An eight-wavelength WDM ring network survivability experiment on NTT's regional fiber network," presented at OFC'97, Dallas, TX, Paper TH05.
- [18] J. Manchester and P. Bonenfant, "Fiber optic network survivability: SONET/Optical protection layer interworking," presented at NFOEC'96, Denver, CO.
- [19] J. Y. Wei, C.-C. Shen, M. J. Post, B. J. Wilson, M. J. Post, and Y. Tsai, "Connection management for multiwavelength optical networking," this issue, pp. 1097-1108.

Mari W. Maeda received the Ph.D. degree in physics from Massachusetts Institute of Technology, Cambridge. Her thesis involved experimental demonstration of squeezed state of light using nonlinearities in atomic vapor.

From 1986 to 1997, she was with Bell Communications Research, Red Bank, NJ, where she worked on optical communications systems research and later developed advanced network management systems for ATM, SONET, and WDM networks. She is currently the Program Manager responsible for Next Generation Internet Program at Defense Advance Research Project Agency, Arlington, VA.